PCT / IB 03 / 0 5 3 2 8
20. 11. 03

Europäisches Patentamt
European Patent Office
Office européen des brevets

Bescheinigung Certificate Attestation

Die angehefteten Unterlagen stimmen mit der ursprünglich eingereichten Fassung der auf dem nächsten Blatt bezeichneten europäischen Patentanmeldung überein.

The attached documents are exact copies of the European patent application described on the following page, as originally filed.

Les documents fixés à cette attestation sont conformes à la version initialement déposée de la demande de brevet européen spécifiée à la page suivante.

Patentanmeldung Nr. Patent application No. Demande de brevet n°

03405393.4

Der Präsident des Europäischen Patentamts;
Im Auftrag

For the President of the European Patent Office

Le Président de l'Office européen des brevets
p.o.

R C van Dijk

Anmeldung Nr:
Application no.: 03405393.4
Demande no:

Anmeldetag:
Date of filing: 30.05.03
Date de dépôt:

Anmelder/Applicant(s)/Demandeur(s):

International Business Machines Corporation
New Orchard Road
Armonk, NY 10504
ETATS-UNIS D'AMERIQUE

Bezeichnung der Erfindung/Title of the invention/Titre de l'invention:
(Falls die Bezeichnung der Erfindung nicht angegeben ist, siehe Beschreibung.
If no title is shown please refer to the description.
Si aucun titre n'est indiqué se referer à la description.)

Detecting netwowrk attacks

In Anspruch genommene Priorität(en) / Priority(ies) claimed /Priorité(s) revendiquée(s)
Staat/Tag/Aktenzeichen/State/Date/File no./Pays/Date/Numéro de dépôt:

Internationale Patentklassifikation/International Patent Classification/
Classification internationale des brevets:

G06F1/00

Am Anmeldetag benannte Vertragstaaten/Contracting states designated at date of filing/Etats contractants désignées lors du dépôt:

AT BE BG CH CY CZ DE DK EE ES FI FR GB GR HU IE IT LU MC NL PT RO SE SI SK TR LI

DETECTING NETWORK ATTACKS

Technical Field

The present invention generally relates to detecting network attacks and particularly relates to methods, apparatus, and computer program elements for detecting attacks on a data communications network

Background of the Invention

The Internet is a wide area data communications network formed from a plurality of interconnected data networks. In operation, the Internet facilitates data communications between a range of remotely situated data processing systems. Such data processing systems each typically comprise a central processing unit (CPU), a memory subsystem, and input/output (I/O) subsystem, and computer program code stored in the memory subsystem for execution by the CPU. Typically, end user data processing systems connected to the Internet are referred to as client data processing systems or simply clients. Similarly, data processing systems hosting web sites and services for access by end users via the Internet are referred to as server data processing systems or simply servers. There is a client-server relationship completed via the Internet between the end user data processing systems and the hosting data processing systems.

The Internet has become an important communications network for facilitating electronically effected commercial interactions between consumers, retailers, and service providers. Access to the Internet is typically provided to such entities via an Internet Service Provider (ISP). Each ISP typically operates an open network to which clients subscribe. Each client is provided with a unique Internet Protocol (IP) address on the network. Similarly, each server on the network is provided with a unique IP address. The network operated by the ISP is connected to the Internet via a dedicated data

processing system usually referred to as a router. In operation, the router directs inbound communication traffic from the Internet to specified IP addresses on the network. Similarly, the router directs outbound communication traffic from the network in the direction of specified IP addresses on the Internet.

A problem faced by many ISPs is the increasing frequency of electronic attacks to the networks they operate. Such attacks include computer virus attacks and so-called "worm" attacks. Attacks of this nature introduce significant performance degradation in networks operated by ISPs. Infected systems connected to the network typically attempt to spread the infection within the network. Many users do not recognize that their systems are infected. It would be desirable to provide technology for triggering disinfection of such systems in the interests of increasing network performance.

Summary of the Invention

In accordance with the present invention, there is now provided a method for detecting attacks on a data communications network having a plurality of addresses for assignment to data processing systems in the network, the method comprising: identifying data traffic on the network originating at any assigned address and addressed to any unassigned address; inspecting any data traffic so identified for data indicative of an attack; and, on detection of data indicative of an attack, generating an alert signal.

The inspecting preferably comprises spoofing replies to requests contained in the data traffic identified. A preferred embodiment of the present invention comprises, on generation of the alert signal, rerouting any data traffic originating at the address assigned to the data processing system originating the data indicative of the attack to a disinfection address on the network. On generation of the alert signal, an alert message may be sent to the disinfection address. The alert

message may comprise data indicative of the attack detected. On receipt of the alert message, a warning message may be sent from the disinfection address to the address assigned to the data processing system originating the data indicative of the attack. The warning message may include program code for eliminating the attack when executed by the data processing system originating the data indicative of the attack.

Viewing the present invention from another aspect, there is now provided apparatus for detecting attacks on a data communications network having a plurality of addresses for assignment to data processing systems in the network, the apparatus comprising: an intrusion detection sensor (IDS) for identifying data traffic on the network originating at any assigned address and addressed to any unassigned address, inspecting any data traffic so identified for data indicative of an attack, and, on detection of data indicative of an attack, generating an alert signal.

The IDS in use preferably inspects the data traffic identified through spoofing replies to requests contained in the data traffic identified. The apparatus may also comprise a router connected to the intrusion detection sensor for rerouting, in response to generation of the alert signal, any data traffic originating at the address assigned to the data processing system originating the data indicative of the attack to a disinfection address on the network. Preferably, the IDS, on generation of the alert signal, sends an alert message to the disinfection address. The alert message preferably comprises data indicative of the attack detected. A preferred embodiment of the present invention further comprises a disinfection server assigned to the disinfection address, the disinfection server sending, on receipt of the alert message, a warning message to the address assigned to the data processing system originating the data indicative of the attack.

The present invention also extends to a data communications network comprising: a plurality of addresses for assignment to

data processing systems in the network; and, apparatus for detecting attacks on the network as herein before described.

The present invention further extends to a computer program element comprising computer program code means which, when loaded in a processor of a data processing system, configures the processor to perform a method for detecting attacks on a data communications network as herein before described.

In a preferred embodiment of the present invention, there is provided a data communications network comprising: a router for connecting a plurality of data processing systems to the Internet; an IDS connected to the router; and a disinfection server also connected to the router. In response to the IDS detecting that one of the data processing systems is infected by an attack, the IDS instructs the router to deflect all network traffic from that attack to the disinfection server. The IDS simultaneously supplies disinfection data to the disinfection server. The disinfection data is indicative of: the nature of the infection; how to disinfect the infecting system; and how to resume normal network connectivity.

There are generally a large number of free IP addresses on a given network. In a particularly preferred embodiment of the present invention, the IDS listens on the network for traffic directed toward the free IP addresses. No such traffic should exist. In the event that a request sent to one of the free IP addresses is detected, the IDS spoofs an answer to the request. The free IP addresses are not in use. Thus, any attempt to contact, for example, a server at such an address is *a priori* suspicious. The IDS then listens for a reply to the spoofed answer. If the IDS detects a diagnosable attack in the reply, it signals the router to divert all traffic from the infected system to the disinfection server. Because, the IDS is interactively spoofing responses to infected systems, it has an accurate view of each attack. Thus, false positives are minimized.

Brief Description of the Figures

Preferred embodiments of the present invention will now be described, by way of example only, with reference to the accompanying drawings, in which:

Figure 1 is a block diagram of a data processing system;

Figure 2 is a block diagram of a data processing network embodying the present invention;

Figure 3 is a block diagram of an intrusion detection sensor embodying the present invention; and,

Figure 4 is a flow diagram associated with the intrusion detection sensor.

Detailed Description

Referring first to Figure 1, a data processing system comprises a CPU 10, an I/O subsystem 20, and a memory subsystem 40, all interconnected by a bus subsystem 30. The memory subsystem 40 may comprise random access memory (RAM), read only memory (ROM), and one or more data storage devices such as hard disk drives, optical disk drives, and the like. The I/O subsystem 20 may comprise: a display; a printer; a keyboard; a pointing device such as a mouse, tracker ball, or the like; and one or more network connections permitting communications between the data processing system and one or more similar systems and/or peripheral devices via a data communications network. The combination of such systems and devices interconnected by such a network may itself form a distributed data processing system. Such distributed systems may be themselves interconnected by additional data communications networks.

In the memory subsystem 40 is stored data 60 and computer program code 50 executable by the CPU 10. The program code 50

includes operating system software 90 and application software 80. The operating system software 90, when executed by the CPU 10, provides a platform on which the application software 80 can be executed.

Referring now to Figure 2, in a preferred embodiment of the present invention, there is provided a data communications network 100 having a plurality of addresses 110 for assignment to data processing systems in the network. In a particularly preferred embodiment of the present invention, the network 100 is in the form of an Internet service installation having a plurality assignable Internet Protocol (IP) addresses 110. The network 100 is connected to the Internet 120 via a router 130. The router 130 may be implemented a data processing system as herein before described with reference to Figure 1 dedicated by appropriate programming to the task to route communication traffic in the form of data packets between the Internet 120 and the network 100 based on IP address data specified in the packets. A first group 140 of the IP addresses 110 on the network 100 are assigned to systems 150 belonging to users of the Internet service. Each system 150 may be a data processing system as herein before described with reference to Figure 1. A second group 160 of the IP addresses 110 on the network 100 are free. More specifically, the second group 160 of IP addresses 110 are not assigned to user systems 150. An intrusion detection sensor (IDS) 170 is also connected to the network 100. The IDS 170 is also connected to the router 130. Details of the IDS 170 will be provided shortly. The router 130 is connected to a disinfection server 180. The disinfection server 180 may be implemented by a data processing system as herein before described with reference to Figure 1.

With reference to Figure 3, in a particularly preferred embodiment of the present invention, the IDS 170 comprises a data processing system as herein before described with reference to Figure 1. The application software 80 of the IDS 170 includes intrusion detection code 200. The data 60 stored

in the memory subsystem 40 of the IDS 170 includes attack identity data 210 and disinfection data 220. The data also includes a record of which of the IP addresses on the network 100 are free and which of the IP of the IP addresses on the network 100 are assigned to data processing systems 100. The record is updated each time another IP address is allocated or an existing IP address allocation is removed. The attack identity data 210 contains data indicative of signatures identifying known attacks. The disinfection data 220 contains data indicative of: the nature of each attack; how to disinfect system infected with each attack; and how to resume normal network connectivity. The attack identity data 210 and disinfection data 220 are cross referenced. The intrusion detection code 200, when executed by the CPU 10, configures the IDS 170 to operate in according the flow diagram shown in Figure 4.

Referring now to Figure 4, in operation, the IDS 170 identifies data traffic on the network originating at any assigned address 140 and addressed to any unassigned address 160. The IDS 170 inspects any data traffic so identified for data indicative of an attack. On detection of data indicative of attack, the IDS 170 generates an alert signal. In a preferred embodiment of the present invention, on generation of the alert signal, any data traffic originating at the address 140 assigned to the data processing system 150 originating the data indicative of the attack is rerouted to a disinfection address on the network 100. In a particularly preferred embodiment of the present invention, the IDS 170 listens on the network 100 for traffic directed toward the free IP addresses 160. Specifically, at block 300, the IDS 170 examines requests sent from addresses 140 on the network 100 to determine, at block 310, if the request specifies one of the free IP addresses 160 as the destination address. If the request does not specify one of the free IP addresses 160, then, at block 320, the IDS 170 waits for the next request to examine. If, however, the request specifies one of the free IP addresses 160, then, at block 330, the IDS 170 spoofs an

answer to the request. The answer is sent to the source IP address on the network 100. The free IP addresses 160 are not in use. Thus, any attempt to contact, for example, a system at such an address is *a priori* suspicious. At block 340, the IDS 170 listens for a reply to the spoofed answer. The IDS 170 may time out if no reply is received within a predetermined period, in which case, at block 320, the IDS 170 waits for the next request to examine. If a reply is however received, then, at block 350, the IDS 170 compares the suspect request and reply with the attack identity data 210 stored in the memory subsystem 40. If, at block 350, the comparison fails to identify an attack, then, at block 320, the IDS 170 waits for the next request to examine. If, however, the comparison at block 350 detects a diagnosable attack in the reply, then the IDS 170 determines that the source system 150 is infected. Accordingly, at block 360, the IDS 170 generates the alert signal. The alert signal is sent to the router 130. The alert signal instructs the router 130 to divert all traffic from the infected system 150 to the disinfection address. Referring back to Figure 1, in a particularly preferred embodiment of the present invention, a disinfection server 180 is located at the disinfection address.

In a preferred embodiment of the present invention, on generation of the alert signal, the IDS 170 sends an alert message to the disinfection address. Preferably, the alert message comprises data indicative of the attack detected. Accordingly, in a particularly preferred embodiment of the present invention, the IDS 170 retrieves the disinfection data 220 corresponding to the attack detected from the memory subsystem 40. At block 370, the IDS 170 sends the alert message containing retrieved disinfection data 370 to the disinfection address at which the disinfection server 180 resides. Then, at block 320, the IDS 170 waits for the next request to examine. Each request, answer, and reply may be embodied in one or more packets of data traffic on the network 100. Accordingly, the signature of each attack may span more than one packet.

In a preferred embodiment of the present invention, the disinfection data 220 sent to the disinfection server 180 contains data indicative of: the nature of the attack detected; how to disinfect the system 150 infected with the attack; and how to resume normal network connectivity. On receipt of the disinfection data 220 from the IDS 170, the disinfection server 180 sets about curing the infected system 150 and restoring the network 100. In another preferred embodiment of the present invention, the disinfection data 220 contains only data indicative of the nature of the attack. The disinfection server then selects, based the nature of the attack, one of a plurality of pre-stored techniques for disinfecting the infected system 150 and/or restoring the network 100 and executes the selected technique. The attacks may take many different forms. Accordingly, the corresponding techniques for disinfection and network restoration may vary widely from one attack to the next.

In a preferred embodiment of the present invention, on receipt the disinfection data, the disinfection server 180 sends a warning message to the infected system 150. The warning message informs the user of the infected system 150 that his or her system 150 is infected. The message may instruct the user to run anti-virus software pre-stored in the infected system 150 to eliminate or otherwise isolate the infection. Alternatively, the message may contain disinfection program code for eliminating the attack from the infected system 150, together with instructions to assist the user in executing the disinfection code on the infected system 150. In another alternative, the message may direct the user to another web site, at which appropriate disinfection program code is provided. In another preferred embodiment of the present invention, the message contains disinfection program code that, when loaded into the infected system, executes automatically, thus eliminating or otherwise isolating the infection in a manner which is transparent to the user. Other disinfection schemes are possible.

In the embodiments of the present invention herein before described, the disinfection server 180 is implemented in a single data processing system such as that herein before described with reference to Figure 1. However, in other embodiments of the present invention, the disinfection server 180 may be implemented by multiple interconnected data processing systems. Such data processing may be distributed or located together in a "farm". Each data processing system in the disinfection server may be dedicated to handling a different attack. The IDS 170 may also be implemented by multiple integrated data processing systems. Alternatively, the IDS 170 and the disinfection server 180 may be integrated in a single data processing system.

The traffic on the network 100 sent from the infected system 150 and deflected by the router 130 to the disinfection server 180 may be logged and/or discarded by the disinfection server 180. In the embodiments of the present invention herein before described, the IDS 170 sends disinfection data to the disinfection server 220. However, in other embodiments of the present invention, once an infection is detected, the IDS 170 may simply instruct the router 130 to deflect traffic from the infected system 150 to the disinfection server 180 without the IDS 170 additionally supplying disinfection data 220 to the disinfection server 180. The disinfection server 180 may then simply act as a repository for traffic originating in the infected system 150, logging and/or discarding traffic it receives from the infected system 150. The logging and discarding may be reported by the disinfection server 180 to an administrator of the network 100. Such reports may be delivered periodically or in real time. The reporting may be performed via, for example, an administration console. However, other reporting techniques, such as printed output for example, are possible. On receipt of such reports, administrators can take actions appropriate for eliminating or otherwise containing the infection of the network 100.

In the embodiments of the present invention herein before described, the IDS 170, router 130, and disinfection server 180 are implemented by data processing systems programmed with appropriate program code. However, it will be appreciated that, in other embodiments of the present invention, one or more of the functions described herein as being implemented in software may be implemented at least partially in hardwired logic circuitry.

It will also be appreciated that the attack detection methods described herein may be implemented by the service provider responsible for the network 100, or at least partially by a third party in the form of a service to the service provider. Such a service may differentiate the service offered by the service provider from the services provided by it competitors. Such differentiated services may be optionally supplied to end users of the network service provided in exchange for an additional premium.

CLAIMS

1. A method for detecting attacks on a data communications network having a plurality of addresses for assignment to data processing systems in the network, the method comprising: identifying data traffic on the network originating at any assigned address and addressed to any unassigned address; inspecting any data traffic so identified for data indicative of an attack; and, on detection of data indicative of an attack, generating an alert signal.

2. A method as claimed in claim 1, wherein the inspecting comprises spoofing replies to requests contained in the data traffic identified.

3. A method as claimed in claim 1, comprising, on generation of the alert signal, rerouting any data traffic originating at the address assigned to the data processing system originating the data indicative of the attack to a disinfection address on the network.

4. A method as claimed in claim 1, comprising, on generation of the alert signal, sending an alert message to the disinfection address.

5. A method as claimed in claim 5, wherein the alert message comprises data indicative of the attack detected.

6. A method as claimed in claim 5, comprising, on receipt of the alert message, sending a warning message from the disinfection address to the address assigned to the data processing system originating the data indicative of the attack.

7. A method as claimed in claim 6, comprising including in the warning message program code for eliminating the attack

when executed by the data processing system originating the data indicative of the attack.

8. Apparatus for detecting attacks on a data communications network having a plurality of addresses for assignment to data processing systems in the network, the apparatus comprising: an intrusion detection sensor for identifying data traffic on the network originating at any assigned address and addressed to any unassigned address, inspecting any data traffic so identified for data indicative of an attack, and, on detection of data indicative of an attack, generating an alert signal.

9. Apparatus as claimed in claim 8, wherein the intrusion detection sensor in use inspects the data traffic identified by spoofing replies to requests contained in the data traffic identified.

10. Apparatus as claimed in claim 8, further comprising a router connected to the intrusion detection sensor for rerouting, in response to generation of the alert signal, any data traffic originating at the address assigned to the data processing system originating the data indicative of the attack to a disinfection address on the network.

11. Apparatus as claimed in claim 8, wherein the intrusion detection sensor, on generation of the alert signal, sends an alert message to the disinfection address.

12. Apparatus as claimed in claim 11, wherein the alert message comprises data indicative of the attack detected.

13. Apparatus as claimed in claim 12, further comprising a disinfection server assigned to the disinfection address, the disinfection server sending, on receipt of the alert message, a warning message to the address assigned to the data processing system originating the data indicative of the attack.

14. Apparatus as claimed in claim 13, wherein the warning message comprises program code for eliminating the attack when executed by the data processing system originating the data indicative of the attack.

15. A data communications network comprising: a plurality of addresses for assignment to data processing systems in the network; and, apparatus for detecting attacks on the network as claimed in any of claims 8 to 14.

16. A computer program element comprising computer program code means which, when loaded in a processor of a data processing system, configures the processor to perform a method for detecting attacks on a data communications network as claimed in any of claims 1 to 8.

ABSTRACT

Described is a technique for detecting attacks on a data communications network having a plurality of addresses for assignment to data processing systems in the network. The technique involves identifying data traffic on the network originating at any assigned address and addressed to any unassigned address. Any data traffic so identified is inspected for data indicative of an attack. On detection of data indicative of an attack, an alert signal is generated.

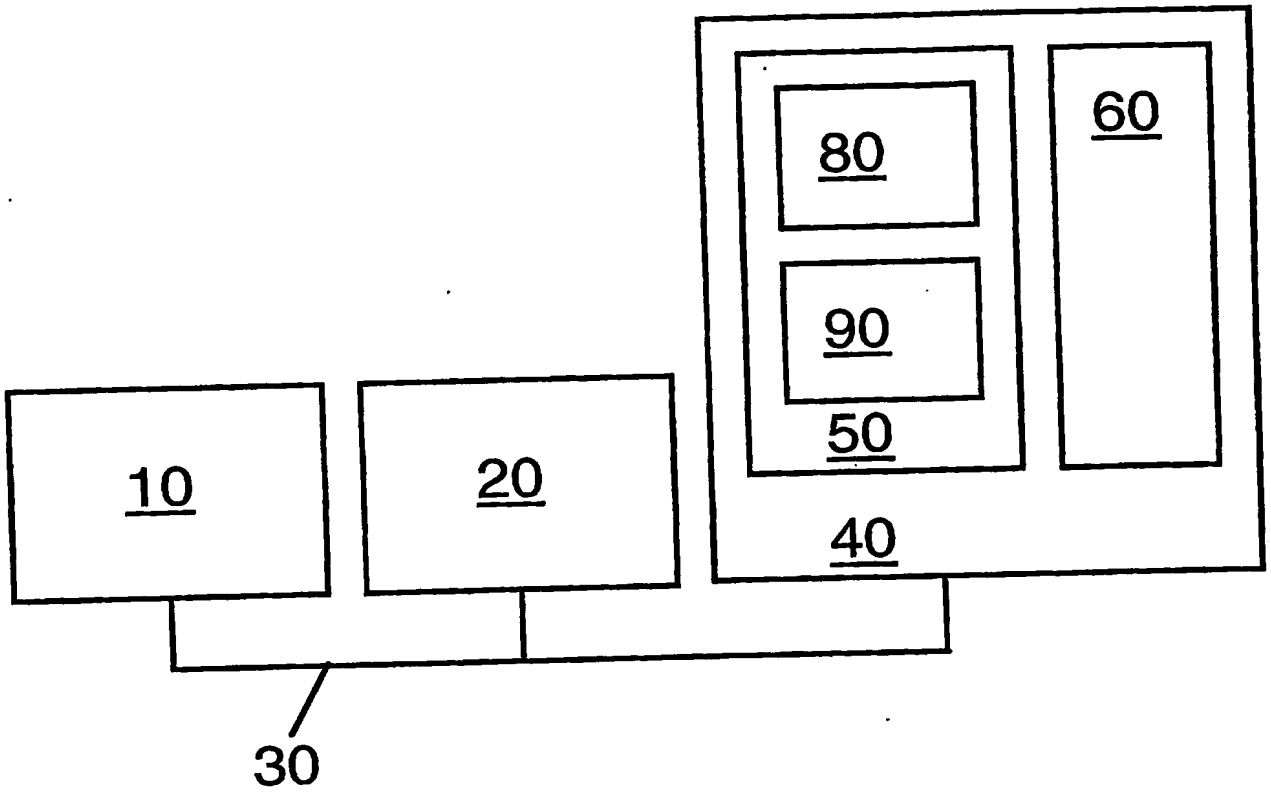

Fig.1

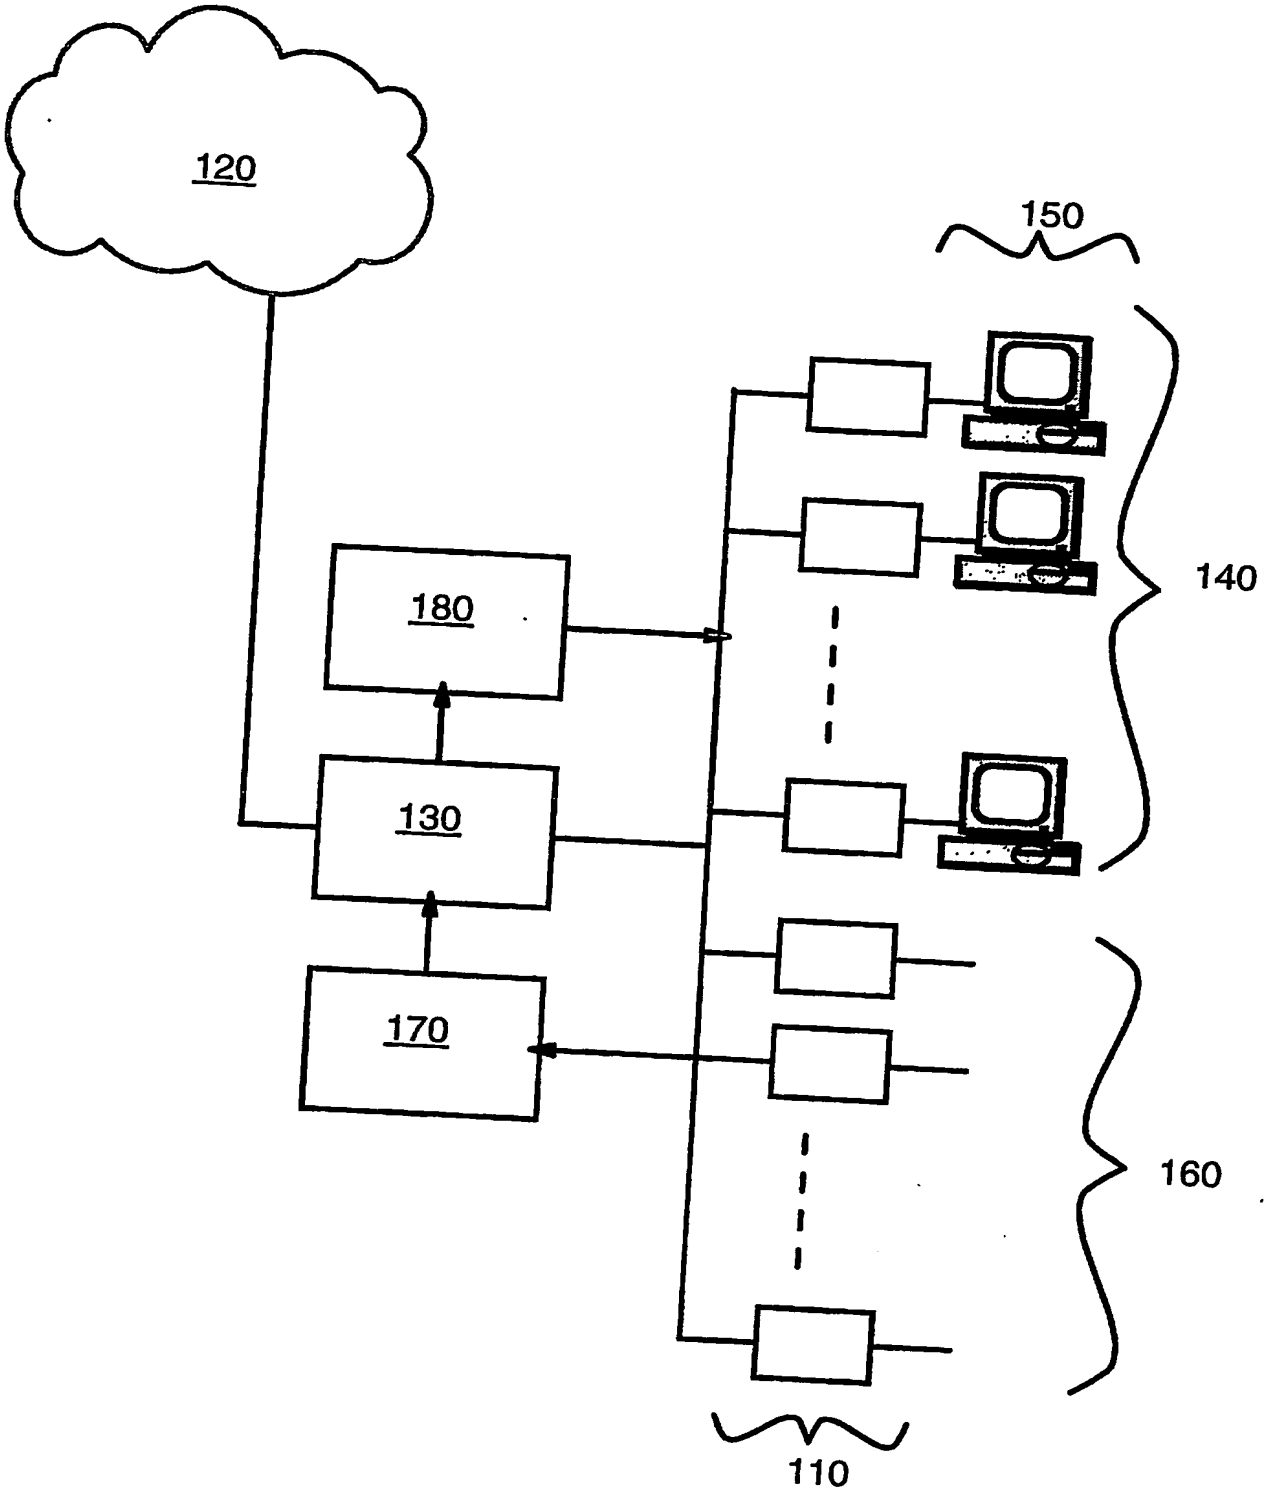

Fig.2

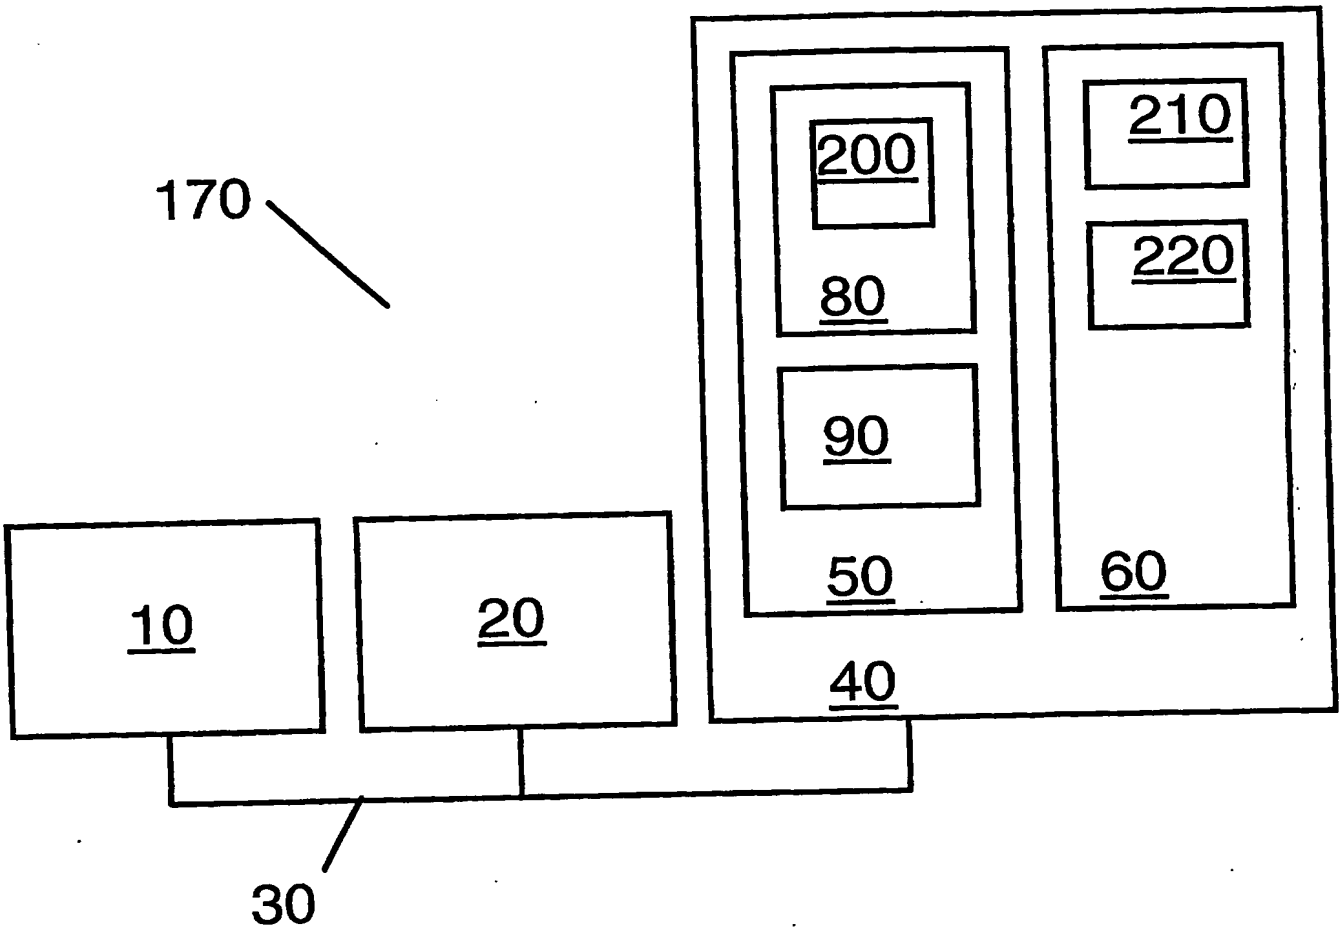

Fig.3

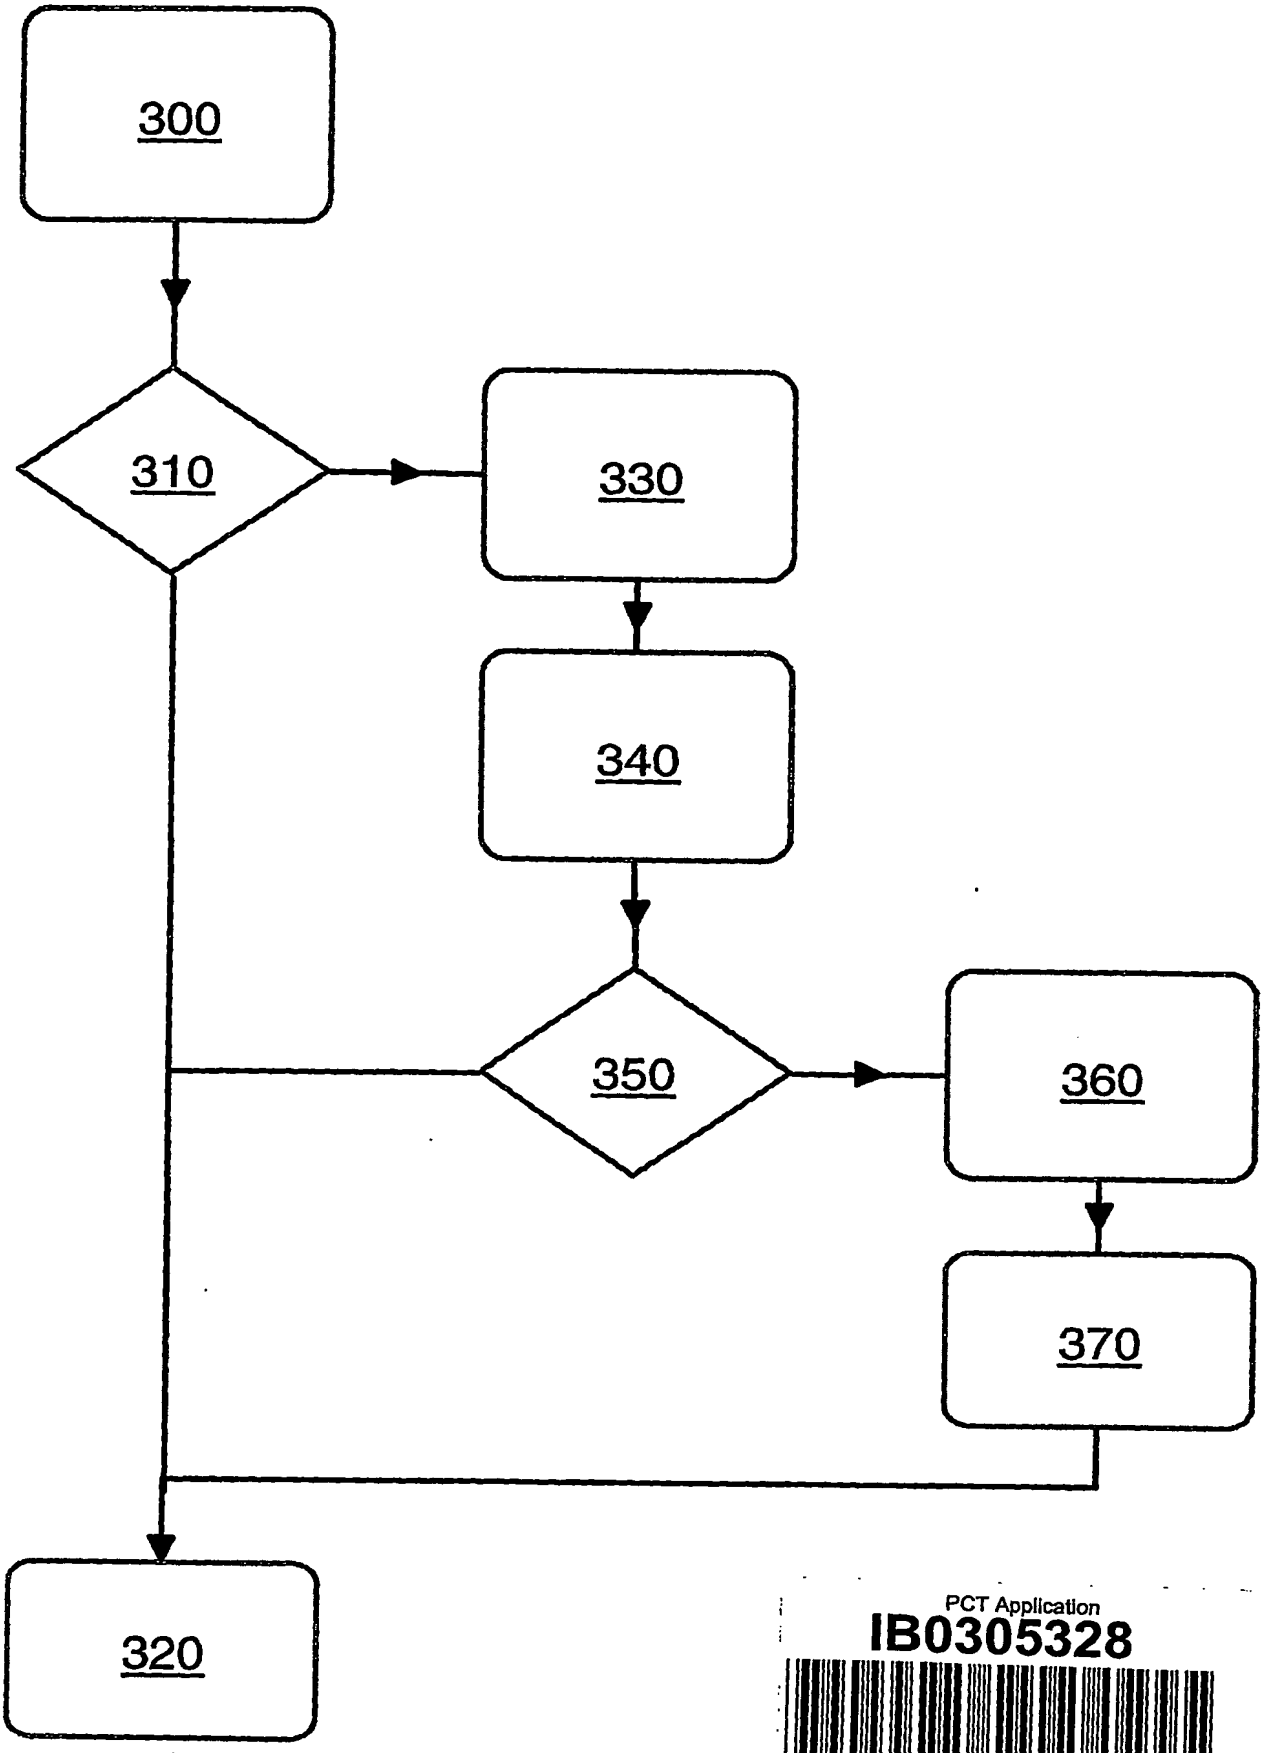

Fig.4